

PCT


WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁶ : G01S 5/16, 17/87, H04N 5/272</p>	<p>A1</p>	<p>(11) International Publication Number: WO 98/54593 (43) International Publication Date: 3 December 1998 (03.12.98)</p>
<p>(21) International Application Number: PCT/GB98/01576 (22) International Filing Date: 29 May 1998 (29.05.98) (30) Priority Data: 9711316.1 30 May 1997 (30.05.97) GB 9711373.2 2 June 1997 (02.06.97) GB (71) Applicant (for all designated States except US): BRITISH BROADCASTING CORPORATION [GB/GB]; Broadcasting House, London W1A 1AA (GB). (72) Inventors; and (75) Inventors/Applicants (for US only): RUSSELL, Richard, Thomas [GB/GB]; 19 Welling Road, Orsett, Essex RM16 3DF (GB). THOMAS, Graham, Alexander [GB/GB]; 'Anderida', Ridge Close, Nutley, Uckfield, East Sussex TN22 3HA (GB). (74) Agent: KAZI, Ilya; Mathys & Squire, 100 Grays Inn Road, London WC1X 8AL (GB).</p>		<p>(81) Designated States: IL, JP, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published With international search report.</p>

(54) Title: POSITION DETERMINATION



(57) Abstract

In the preferred implementation, the position of an object, for example a studio camera is determined by means of a camera which views several markers disposed about a studio ceiling, the markers being patterned as a series of light and dark rings to encode information

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

POSITION DETERMINATION

The present invention relates to position determination and is particularly, but not exclusively, concerned with derivation of camera position for so-called "virtual studio" television production. In such applications, a virtual scene, for example a computer-generated background, is superimposed onto a real scene, for example actors. It is important to know the exact position and orientation of the television camera in the studio, so that the virtual scene can be computed to correspond to the viewpoint of the real camera, to ensure correct registration between real and virtual elements in the scene.

There are a number of commercially-available systems that can provide camera position information. Most are based on mechanical camera mountings, such as robotic pedestals, pedestals mounted on tracks, or robot arms, which have sensors to measure position and orientation. These systems cannot be used with hand-held cameras, and can be bulky and difficult to use.

There are also methods that work without mechanical sensors, but instead use a patterned blue background, which is visible in the camera image. By analysing the video signal, these methods can deduce the orientation and position of the camera. One example of such a method is described in our earlier GB-A-2271241, which derives pan, tilt and zoom using an arbitrary patterned background. A further example using the same technique is described in WO-A-95/30312, which uses a particular type of pattern to enable pan, tilt, zoom and position to be determined. However, these methods rely on the presence of a two-tone blue background. This is inappropriate in some situations, for example when it is desired to extract shadows of objects from the foreground image. Also, in some situations there may be little or no blue background visible, for example during a close-up shot of an actor. Sometimes it may also be required to place a virtual object against a real background, in which case there will be no blue background at all.

A method for locating the position of a head-mounted display that has been used in the field of Augmented Reality is described in the paper by Azuma et al entitled "A Demonstrated Optical Tracker with Scalable Work Area for Head-Mounted Display Systems" published in ACM Computer Graphics: Proceedings of the 1992 Symposium on Interactive 3D Graphics (Cambridge, Massachusetts, April 1992), pp. 43-52. This method uses a number of infra-red LEDs mounted in ceiling panels, viewed by four upward-looking sensors mounted on the user's headset. The sensors each provide information of the co-ordinates of bright points in the image (the LEDs) and from these co-ordinates, the known position of the LEDs and the geometry of the sensors, the position of the headset is computed. The inventors contemplated applying Azuma's method to the problem of determination of camera position. However, Azuma's method is not intended to be used for determining the position of a camera in a large studio, but is designed instead to work in a relatively small volume, where the total number of LEDs is small. The inventors identified several potential problems in applying such a technique to the field of camera position determination.

In a television studio, the camera could move many tens of metres, so the method of Azuma et al would require a very large number of ceiling markers of which only a small proportion would be visible at any one time. Azuma's method relies on active control of the LEDs to identify the LEDs. The control electronics and wiring required to implement this in a large studio would be complex and impracticable. Furthermore, the ceiling of a television studio generally contains a number of lights and other objects that could make identifying the markers much more difficult than in the carefully controlled conditions in which Azuma's method is designed to function; the bright spot sensors used by Azuma would generate spurious signals. A still further problem the inventors have identified is that the set of markers that the camera can see will change, not only due to markers coming into and out of the field of view as the camera moves, but also due to markers being obscured by objects such as microphone booms and studio lights. Using Azuma's method,

each time a marker appears or disappears, there is likely to be a small but sudden change in the computed position and orientation of the camera, since any errors in the camera calibration or the measurement of the marker positions will lead to results that depend on which markers are being used.

5

Thus, Azuma's method cannot be directly applied to the present problem. The inventors have developed novel techniques which are particularly suited to the exacting requirements of determining camera position in a television studio, which overcome or alleviate the drawbacks of conventional techniques. References in this specification to a video camera are intended to include any camera capable of obtaining video images; although the video camera is advantageously a conventional television studio camera, no limitation to a particular type of camera or intended use is implied.

15

In a first aspect, the present invention provides a method of determining the position of an object comprising:

providing a plurality of markers at respective reference positions, at least some of the markers being patterned to encode identification information, the pattern providing a reference point for each marker;

20

storing information including a measure of the positions of the markers and information identifying the patterned markers;

obtaining an image of at least a sub-set of said plurality of markers from a camera associated with the object;

25

processing the image to identify the positions of said markers in the image and, for each patterned marker in the image, decoding said identification information;

determining a measure of the position of the object based on said processing and decoding and based on said stored information.

30

Thus, with the invention, at least some, and preferably all, markers are patterned to encode information identifying the marker. In this way, it becomes possible to determine not only relative movements of the object but also the

absolute position by "looking up" the positions of each patterned marker. This may enable much greater freedom of movement of the object over a wider area whilst allowing absolute position to be measured accurately. In addition, determination of relative movement may be simplified or improved.

5

Preferably, the pattern comprises concentric shapes, preferably substantially closed rings; the use of concentric rings to encode the information facilitates location of the markers, as the centre of the rings can conveniently be located and provides a convenient reference position. Preferably, identifying the positions of the markers includes identifying the centres of the concentric rings as a reference point for each patterned marker. Although the rings are most preferably concentric to within the resolution of the camera means in use as this greatly facilitates identification and decoding, they need not be exactly concentric; in such a case, eccentricity may be used to provide a measure of angular orientation, and the reference point may be determined based on, for example, the centre of the innermost ring.

15

Preferably, the patterned markers encode information as a series of light and dark regions. This enables a monochrome camera to be used to identify the markers, allowing higher resolution to be obtained at lower cost. The camera may operate in the visible region of the spectrum, the patterned markers having visible markings; this facilitates identification of the markers during setting up by a user. However, infra-red (or ultra-violet) light may be used.

20

25

Preferably the identification information is encoded in binary form as the brightness or reflectivity of each ring. This can provide a robust and simple coding scheme. The rings in such a scheme are preferably of known inner and outer diameter and preferably of substantially equal thickness. If the dimensions or relative dimensions of each ring are known, it is not necessary for the rings to be visibly delimited; adjacent rings can be contiguous even if of the same shade. In addition, a measure of the distance to each marker can be

30

derived based on the size of the marker in the image.

Alternatively, the information can be encoded in binary form as the thickness of each ring, successive rings preferably alternating between light and dark.

The rings are preferably substantially circular (this facilitates detection and decoding) and preferably all of the same or similar shapes. However, other shapes (for example squares, rectangles or polygons) may be used; the term "ring" as used in this specification is intended to encompass such non-circular rings.

Surprisingly it has been found that good results can be obtained with monochrome patterns and binary encoding, even though this may require several rings. This has been found to be due in part to the better resolution and linearity generally attainable with monochrome cameras than with colour cameras. In addition, monochrome encoding enables a much higher dynamic range between light and dark to be attained and provides lower sensitivity to changes in ambient lighting.

Indeed, this binary encoding enabling monochrome detection is an important feature which may be provided independently in a second aspect, in which the invention provides a method of determining the position of an object comprising:

providing a plurality of markers at respective reference positions, at least some of the markers being patterned to encode identification information in binary form as a series of light and dark regions, the marker also providing a detectable reference point;

storing information including a measure of the positions of the markers and information identifying the patterned markers;

obtaining an image of at least a sub-set of said plurality of markers from a camera, preferably monochrome camera means, associated with the object;

processing the image to identify the positions of said markers in the image and, for each patterned marker in the image, decoding said identification information;

5 determining a measure of the position of the object based on said processing and decoding and based on said stored information.

Surprisingly, if binary encoding is used, the markers can be relatively complex and still be reliably decoded. Thus, in preferred arrangements, the markers encode at least 3, and more preferably 4 or more bits of identification
10 information. This facilitates identification of the markers, as a larger number of markers can be uniquely identified. Preferably, at least one guard ring or error-correcting ring is provided, so the markers contain at least 4 or more rings.

In a related third aspect, the invention provides a set of patterned
15 markers each comprising a series of concentric rings, each marker encoding a plurality of bits of information in binary form as the brightness or reflectivity and/or thickness of each ring, the information encoded varying between the markers. Some markers may be repeated (i.e. encode the same information) within the set.

20 The invention also provides use of a series of concentric rings to encode a plurality of bits of information in binary form as the brightness or reflectivity and/or thickness of each ring.

25 The invention further provides, in a fourth aspect, a method of producing a marker encoding a plurality of bits of information, the method comprising obtaining the information to be encoded in binary form and providing on a backing a series of concentric rings, each ring having either a high or a low brightness or reflectivity, the thickness and brightness or reflectivity of each ring
30 being selected according to a pre-determined coding scheme to encode the information.

This may be implemented as a method of producing a marker encoding a plurality of bits of information, the method comprising obtaining the information to be encoded in binary form and providing on a backing a series of concentric rings, each ring corresponding to a bit of information to be encoded and having a first brightness or reflectivity to encode a binary "0" and a second brightness or reflectivity to encode a binary "1".

A variety of coding schemes may be used, including self-clocking and self error-correcting schemes. For example, known linear barcoding schemes may be used with successive bars corresponding to successive rings.

Although a number of coding schemes may be used, it is found that, surprisingly, a simple coding scheme in which the rings are of substantially constant thickness and each have a brightness or reflectivity corresponding directly to a bit of information can provide a compact marker which can be reliably and readily decoded.

One or more parity bits or error-correcting bits may be included (preferably one parity bit for the "even" rings and one for the "odd" rings); this enables offset errors to be detected.

Decoding the markers may include comparing the identification information to predicted identification information based on marker positions and correcting or verifying the identification information based on the results of the comparison. For example, if the decoded marker identification information differs from the predicted information by an amount corresponding to a single bit or a shift of all bits, this may be interpreted as a corresponding read error and corrected "intelligently".

The backing preferably has a surface coating of said first or second reflectivity or brightness (for example white or retro-reflective or black); in this way, only rings of the other reflectivity or brightness need be actively provided,

for example by applying retro-reflective material or pigment.

5 The dark regions are preferably substantially black. The light regions may be white, but may be advantageously formed by retro-reflective material, to be illuminated by a light substantially optically coincident with the camera axis. This enables a high level of contrast to be attained, and facilitates detection of markers. The use of retro-reflective material and a light source associated with the camera renders detection of the markers less sensitive to changes in ambient lighting conditions, as may occur in a television studio. In
10 place of retro-reflective material, fluorescent pigments may be employed to enhance the contrast.

Preferably each patterned marker has an outer region, preferably a relatively wide annular region, which is substantially the same shade for each
15 marker, preferably dark; this may facilitate distinguishing the marker from other items, for example studio lights. Preferably the ring immediately inside the outer region is of the opposite shade to said outer region, preferably light; this may facilitate determination of the size of the pattern and enable more reliable decoding of the information. The outer region may contain one or more
20 markings enabling the angular orientation of the marker to be determined.

The innermost ring or rings may be of predetermined brightness to facilitate detection of the centre of the marker, or may be used to encode identification information.

25

As an alternative to binary encoding, the rings may vary in colour, each ring encoding more than one bit of information. For example, a range of colours may be employed, with a particular bit pattern assigned to a selected colour, preferably based on an RGB encoding scheme. This may allow more
30 information to be encoded in a smaller number of circles (for example, if 8 secondary colours are used, each ring can encode 3 bits).

Preferably, at least some markers encode identification information uniquely identifying that marker among the plurality of markers. This enables the absolute position to be determined without needing to know an initial starting position. The markers are preferably so disposed that the position of the object can be uniquely determined at all points within a given volume.

The method may include adjusting the determined position to smooth out sudden changes in determined position as markers are revealed or obscured. This may alleviate the problem of small, but nonetheless sudden and very visible changes in determined object position due to small errors in determination of position of the markers as different markers move in and out of the field of view of the camera. Adjusting is preferably achieved by applying correction factors to the stored or measured positions of the markers so that determination of object position is based on corrected marker positions, the correction factors being calculated so that the corrected marker positions tend to mutually self-consistent values, the rate of variation of the correction factors being kept below a pre-determined level.

Adjustment may be effected by tracking movement of the markers as the object position changes and generating 3-dimensional correction vectors. In this way, stable, refined marker positions may be determined. However, tracking markers over several frames may be computationally intensive, and will fail to produce stable values for marker positions where errors are attributable to any non-linearity in the camera or lens. Thus, a preferred, simpler implementation of adjustment comprises generating correction factors with 2 degrees of freedom (for example displacement vectors parallel to the image plane or parallel to the reference surface) and correcting the marker positions to tend to render the marker positions self-consistent for an image frame or for a series of image frames.

The rate of variation may be limited so that adjustment of the correction factors occurs over a period of a few seconds (preferably at least about 2

seconds, preferably less than about 15 seconds, typically about 5–10 seconds), so that a gradual drift in the determined position occurs which is much less perceptible than a sudden shift.

5 A measure of the error or accuracy of position determination may be provided based on the self-consistency of determined position calculated for each marker. This may be based on the magnitude of said correction factors.

10 In a fifth aspect, the invention provides a method of determining the position of an object based on identification of a plurality of markers at respective reference positions, the method comprising applying correction factors to stored measures of the marker positions to produce respective corrected marker positions and determining the object position based on the
15 corrected marker positions, wherein the correction factors are calculated so that the corrected marker positions tend to mutually self-consistent values, the rate of variation of the correction factors being limited to a pre-determined level.

20 Although the markers may be adhered to a flat reference surface, the markers are preferably positioned at varying distances from a reference surface, preferably the ceiling of a room in which the object is moved; this is found to provide a surprising improvement in the accuracy of detection of movement of the object. Particularly in such a case, the stored information preferably contains a measure of the distance of each marker from the
25 reference surface. This may be achieved by storing the three-dimensional co-ordinates of each marker with respect to a defined origin (in cartesian or any polar or curvilinear form which is convenient for the position determination algorithm).

30 In the preferred application, the object is a video camera, and the camera means comprises a separate, generally upward pointing, camera mounted on the video camera.

In a sixth, apparatus, aspect, the invention provides apparatus for position determination comprising a camera for mounting on an object whose position is to be determined; memory means arranged to store measures of the positions of a plurality of markers and to store information identifying patterned encoded markers; image processing means arranged to process an image output by the camera to identify the positions of markers in the image; decoding means for decoding information encoded in the patterned markers; and position determining means arranged to determine the position of the object based on the output of the image processing means, the decoding means and the information stored in the memory means.

The memory means, image processing means and position determining means may all be integrated into a single computer. However, preferably, at least part of the function of the image processing means is provided by a hardware accelerator arranged to identify one or more markers or to decode one or more patterned markers or both.

The apparatus preferably further includes means for applying correction factors to stored measures of the marker positions to produce respective corrected marker positions and determining the object position based on the corrected marker positions, wherein the correction factors are calculated so that the corrected marker positions tend to mutually self-consistent values, the rate of variation of the correction factors being limited to a pre-determined level.

In cases where the position of a video camera (separate from said camera means) is determined, position determination may be supplemented by information obtained from the video camera image. In addition, information concerning the video camera zoom and focus settings may be input. One or more encoded markers may be placed in the field of view of the video camera. Provision of encoded markers at specific positions in the field of view is much easier to implement in practice than provision of a patterned background for an

entire studio, as used in prior art methods. Moreover, the encoded markers may enable more accurate measurement than a conventional background in which it is difficult to provide accurate patterning over a large area.

5 Thus, the method of the first or second aspect may further comprise providing at least one supplementary marker at a given position so as to be viewable by the video camera, the or each marker having identification information encoded thereon in binary form as a series of concentric rings of two tones, the tones being selected to enable the marker to be keyed out of the
10 video camera image by chroma-keying, said determining the position of the video camera being based on the position of the supplementary marker, if present, in the video camera image.

 This supplementary positioning may be provided independently, as a
15 seventh aspect, in a method of determining or correcting a measure of the position of a video camera, the method comprising providing at least one marker at a given position viewable by the camera, the or each marker having identification information encoded thereon in binary form as a series of concentric rings of two tones, the tones being selected to enable the marker to
20 be keyed out of the video camera image by chroma-keying, a measure of the position of the video camera being determined based on the position of the marker, if present, in the video camera image. Complete position determination may be effected by use of several such markers so that preferably at least three individual markers are always visible, or by additionally using mechanical
25 sensing, or other optical methods, with the encoded markers providing a check or reference point. Preferably a set of preferably at least three individual markers are mounted on a substantially rigid support to form a cluster of markers, the markers preferably being non co-planar.

30 In an eighth aspect, the invention provides marker apparatus for positioning in the field of view of a camera for use in determining the relative position of the camera and the marker apparatus, the apparatus comprising at

least three non-co-planar patterned markers mounted at predetermined relative positions on a substantially rigid base, the patterned markers each having identification information encoded thereon in binary form. The information is preferably encoded as a series of concentric rings, preferably of
5 two tones, the tones being selected to enable the marker to be keyed out of the video camera image by chroma-keying.

Reference has been made above to a camera means mounted on the object and markers disposed around a volume in which the object is located.
10 This is indeed preferable for determining the position of a camera in a large studio. However, it will be appreciated that the technique may be employed to determine the position of one or more objects on which markers are located using one or more cameras mounted in the volume in which the object is located. For example, the positions of several cameras within a relatively small
15 studio may be determined by a single camera mounted in the ceiling of the studio and having a sufficient field of view to view all cameras; this may lead to a saving in cost.

Thus, in a ninth and final aspect, the invention provides a method of
20 determining the relative positions of camera means and an object having a plurality of markers, preferably at least three, mounted thereon, the relative positions of the markers and the object preferably being substantially fixed, each marker being patterned to encode identification information in binary form, preferably as series of concentric rings, the method comprising storing the
25 relative positions of the markers and information enabling the identification information to be decoded, identifying the positions of the markers and decoding the identification information thereon, and, based on the stored information, determining a measure of the relative positions and orientations of the camera means and the object. The method is preferably employed to
30 determine the positions of several objects relative to the camera means, the objects being distinguished based on the identification information. The camera means may comprise several cameras having different viewpoints.

An embodiment of the invention will now be described in more detail, with reference to the accompanying figures in which:

Figure 1 shows the arrangement in a studio;

Figure 2 shows an example of a marker;

Figure 3 shows a block diagram of an algorithm that may be used to locate markers in the image;

Figure 4 shows the aperture of the filters used in the marker detection system;

Figure 5 shows in detail the process of counting the number of samples above a given threshold in a rectangular area of the image;

Figure 6 shows the method used to prevent the appearance or disappearance of markers from causing a sudden change to the computed camera position;

Figure 7 depicts a marker suitable for placing in the field of view of a video camera.

The way in which a camera position system according to the present invention may be applied in a virtual studio is shown in Fig. 1. A number of markers 1 are fixed on the studio ceiling. Further markers are fixed a short distance below the ceiling, mounted on rigid poles 2. Although the system can work with markers in a plane, greater accuracy may be achieved with markers arranged at different heights, as shown. The system can alternatively be implemented with markers placed in positions other than on the ceiling (such as on a wall, or even on the floor), although mounting them on the ceiling is likely to be most practical in many situations.

The markers are viewed by an upward-looking camera 3 mounted on the side of the normal studio camera 4. The upward-looking camera, referred to as the auxiliary camera, is preferably a monochrome camera, and is preferably progressively-scanned. The markers may be illuminated by a light 12. The signal from the camera is analysed by an analyser 5 to compute the camera position and orientation of the auxiliary camera. The analyser then computes the position and orientation 6 of the studio camera, by compensating for the relative orientation and position of the auxiliary camera 3 and the studio camera 4.

10

The analyser can be implemented using either special-purpose hardware, or by a combination of a frame grabber and a general-purpose computer. Information on the settings of zoom and focus of the studio camera may be gathered using conventional sensors 7, and delayed by a delay 8 to compensate for the delay in the processing 5. The renderer 9 computes the view of the virtual world according to the measured position, orientation and zoom values. Such renderers are available commercially, and are usually based on graphics computers. The normal video signal from the camera is delayed by delay 10 to compensate for the delays in the camera position calculation and the image rendering, before being combined with the rendered image using the conventional keyer 11.

15

20

An example of a marker according to the present invention is shown in Figure 2. The marker shown has bright concentric rings against a dark background. The bright parts are preferably made using retro-reflective material, illuminated by a light adjacent to the camera viewing the markers as shown in Fig. 1. A suitable form of retro-reflective material is 'Reflective silver adhesive transfer film', part no. 8850, made by 3M. In the example shown in Fig. 2, the marker is divided into a number of concentric rings of equal width. The outer ring is always white, and serves to delimit the edge of the marker. Each inner ring is either bright or dark, to indicate a '1' or '0' in the corresponding bit of the marker number. Other types of barcode, for example

25

30

based on the spacing between transitions rather than the brightness of particular bands, are also possible. The example in Fig. 2 has 6 inner rings, giving a 6-bit code that can represent numbers in the range 0-63. It is preferable not to use all possible code numbers, thereby allowing some redundancy. For example, one bit could be used as a parity bit, so that an error in reading one bit can be detected. It is also preferable not to use the code whose representation is a solid white circle (63 in this case), since other objects, such as a studio light against a black background, could be incorrectly interpreted as a marker.

Positioning of the markers for a typical application will now be described. For studios with floor areas up to about 50 square metres, a 9-bit barcode is sufficient to uniquely identify every marker required. For larger studios, some marker numbers can be re-used, and the system can deduce which marker is which by reference, for example, to the nearest uniquely-numbered marker. The minimum size of the markers must be chosen to ensure that the barcode can be successfully read by the auxiliary camera. The present barcode design and reading algorithm preferably requires a minimum spacing between rings corresponding to about 1.5 pixels on the camera sensor to allow reliable reading. For example, if the markers were about 4m above the auxiliary camera, each contained a 9-bit barcode, and a conventional-resolution camera with a minimum field of view of 30 degrees was used, the minimum marker diameter is approximately 12cm. The markers should be positioned so that the auxiliary camera can view a minimum of three, and ideally 10 or more, at any one time. Allowance must be made for some markers being obscured by lights, microphone booms, and so on. For example, in a situation with markers at heights of 4m and 4.5m, a maximum working height of the camera of 2m and the minimum field of view of the auxiliary camera of 30 degrees, positioning markers in a square lattice at 0.4m intervals would ensure that at least 12 were potentially visible, which should allow good performance even if half of these were obscured by obstacles such as lights. Although the markers can be positioned in a coplanar manner, the accuracy of the computed camera position

is significantly improved if the markers are placed at differing distances from the camera. For example, some markers could be mounted directly on the ceiling of a studio, with others mounted on girders or poles hanging a little way below. The position of each marker should ideally be known to an accuracy of approximately 1mm in order to meet the overall accuracy requirements. This could be achieved by measuring the position of the markers using a theodolite, although this would be time-consuming. An alternative method has been developed. First, the positions of the markers are measured roughly, to an accuracy of a few cm. A number of images of the markers are then analyzed to yield accurate positions.

The first task that analyser 5 in Fig. 1 carries out is to identify possible markers in the image. The marker design shown in Fig. 2 allows this to be achieved with a simple algorithm that requires relatively little computation. A block diagram of such an algorithm is shown in Fig. 3. It is based on the use of filters with rectangular apertures, which can be implemented using one-dimensional horizontal and vertical filters. Note that for clarity, Fig. 3 does not show any additional delays to compensate for the differences in processing delay of the various elements. An example of the apertures used is shown in Fig. 4. The inner rectangle A is chosen to be large enough to completely cover all rings of the marker, for the largest expected size of a marker in the image. The outer rectangle B is chosen to fit within the black outer surround of the marker, for the smallest expected size of a marker in the image.

The operation of the marker detector shown in Fig. 3 will now be described in detail. The video signal from the upward-looking camera is input to a sample counter 20 which counts the number of samples above a threshold T1 in a rectangle A centred on the current pixel. This threshold is chosen to correspond to a brightness level between the values of black and white as they appear on the target. This can be implemented using the circuitry shown in Fig. 5, which will now be described. A comparator 40 compares the input signal 39 with a threshold T, and produces a signal 41 which is unity for

samples whose values are greater than or equal to T , and zero for samples less than T . This signal is fed to a vertical summer 46, which computes the sum of its input samples over a vertical aperture of height ' h '. This signal is then passed to a horizontal summer 47, which performs a similar operation for a horizontal aperture of width ' w '. Summers such as these are used in many image processing operations, but for completeness their operation will be described. The following description relates to the vertical summer; the horizontal summer operates in a similar manner. The delay 42 delays the signal by the number of lines equal to the height of the rectangle A. The subtractor 43 forms the difference between the input and output of the delay. The line delay 44 and adder 45 form an integrator which accumulates the output of the subtractor 43. As samples enter the region of height h lines, they are therefore accumulated in the integrator, and as they leave the region, they are subtracted from the accumulated value. The output of adder 45 is thus equal to the number of samples above threshold T in a column of pixels of height h lines. Similarly, the output of the horizontal summer 48 is equal to the number of samples whose values are equal to or greater than the threshold T over a rectangular area of h lines by w pixels.

The output of the sample counter 20 in Fig. 3 is fed to a comparator 24. The comparator produces a 'true' output whenever there are more than N_1 samples in rectangle A over threshold T_1 . The values of T_1 and N_1 are chosen so that this condition is satisfied whenever a marker lies within the rectangle A. Further similar processing is used to check whether the marker is surrounded by a black border. Sample counter 21 counts the number of samples above a threshold T_2 in rectangle A. T_2 is lower than T_1 , and is chosen to be just above the nominal brightness of the black background. Sample counter 22 performs the same task for rectangle B. The difference between the two values corresponds to the number of samples above the threshold T_2 in the region between rectangles A and B in Fig. 4. The comparator 25 tests for this number being less than N_2 , which should be a small number compared to the number of pixels in this region. Its output is

'true' whenever this condition is satisfied.

The possible presence of a marker is indicated when the outputs of comparators 24 and 25 are both 'true'. However, this will probably occur for many adjacent sample positions, centred around the actual marker centre. To identify the exact position of the marker centre, some further processing is carried out. The image signal is first clipped between two levels T3 and T4 by the clipper 23. That is, if the input to the clipper is between T3 and T4, the output is equal to the input. If the input is less than T3, the output is T3, and if the input is greater than T4, the output is T4. The levels T3 and T4 are chosen to ensure that the black surround of the markers lie below T3 and the white level of the rings lies above T4. This prevents any slight changes in illumination level across the marker from disturbing the measurement of the centre position. The clipped signal is then processed to estimate the position of the centre of the marker. The processing is carried out separately for the horizontal and vertical position. To estimate the horizontal position, the values of the clipped image signal are summed independently in the left and right-hand halves of rectangle A. This may be carried out with an arrangement similar to that shown in Fig. 5, without the initial comparator 40. The sum of the values in the right-hand half is subtracted from those on the left. This is carried out by process 26 in Fig. 3. This value will be positive when the marker appears predominantly in the right-hand half, and conversely will be negative when it lies predominantly in the left. This signal is examined for its sign changing from positive to negative on successive pixels by the sign-change detector 29, indicating that the position of rectangle A has moved from being just to the left to just to the right of the centre of the marker. This detector produces a 'true' output when this occurs. The horizontal position of the marker centre is estimated to sub-pixel accuracy by interpolator 31 which estimates the position at which the output of 26 passed through zero. This can be achieved, for example, by linear interpolation between the two values either side of zero. Similar processing (27, 30) is used to locate the vertical position of the marker centre.

This method of finding the centre of a marker has the advantage of being computationally efficient, since the majority of the processing requires only a small number of adders, subtractors and delay elements. However, more sophisticated methods of finding the centre of the image of a circle may also be used. Many possible methods are known from the literature, based for example on computing the centre of gravity, or tracing the edge contour. Such methods could be applied to this invention, and may offer the possibility of obtaining greater accuracy at the expense of higher computational requirements.

10

The presence of a marker is flagged whenever the outputs of 24, 25, 29 and 30 are all 'true'. When this happens, the position of the marker, interpolated to sub-pixel accuracy, may be read from interpolator 31.

15

20

25

30

Once the centre of a marker has been found, the image around this position may be processed to read the marker number represented by the barcode. The algorithm used for performing this will depend on the type of barcode used. For example, with the type of barcode shown in Fig. 2, one way to proceed is as follows. A number of consecutive video samples are selected, starting at the pixel nearest the marker centre and extending to the edge of rectangle A. These pixels will represent samples along a radius of the marker. The outer edge of the marker is located by examining the samples, starting from the outermost one and working inwards, looking for a sample above a given threshold. The threshold would normally be chosen to be mid-way between the video levels of white and black. Once the edge is located, the positions along the radius corresponding to the centres of each bit of the barcode may be examined, and compared to the threshold. This process may make use of interpolation to estimate the video levels at locations between existing samples. For each bit position, the bit in the code is set to 1 if the video level is above the threshold, else it is set to zero. This process may be repeated for several sets of samples starting at the centre, for example working to the left, the right, above, and below the centre. This allows a number of independent readings

of the barcode to be made. The readings may be compared, and the marker either accepted or rejected. For example, a marker could be accepted as valid if three out of four readings of the barcode agree, or alternatively a unanimous verdict may be required. Further modifications are possible, for example
5 changing the thresholds used in the barcode-reading process to obtain the best possible agreement between readings along different radii.

If the barcode-reading process produces an acceptable result, a database listing all markers in the studio is examined to check that the marker exists, and
10 if so to determine its coordinates in the studio. If the marker is not in the table, it may be discarded as corresponding to an erroneous reading. Alternatively, some error-correction techniques could be employed to attempt to correct for mis-read bits in the barcode. There are several well-known error correction methods that can be employed, at the expense of incorporating redundancy in
15 the barcode. This requires the addition of more bits to the code, requiring larger markers. Other information could be used to help resolve an unsuccessfully-read barcode, such as an estimate of which marker might be expected at this position in the image, given knowledge of the camera position measured on the previous frame.

20

From the list of marker positions in the image and the corresponding positions of the markers in the studio, it is possible to compute the camera position and orientation. For example, this may be carried out using the standard photogrammetric technique of 'bundles adjustment', which computes the
25 camera position and orientation that minimise the sum of the squares of the distances between the positions in the image at which the markers would be expected to appear (given the known position of the markers in the studio and the internal camera parameters) and where they are actually found. Suitable techniques for performing bundles adjustment are described in "Manual of
30 Photogrammetry, 4th Edition" American Society of Photogrammetry, Falls Church, VA., 1980, edited by Slama, C. C. . More sophisticated methods may alternatively be used that do not require complete knowledge of the internal

camera parameters or marker positions.

However, performing this computation using the marker positions exactly as found in each image can result in a sudden change in the computed camera position as different markers come into and out of view, as explained earlier. It has been found to be much more acceptable for the computed camera position to change slowly in such circumstances. This is particularly helpful in a situation where a marker is intermittently detected, since the effect of it disappearing for one frame and then reappearing can be significantly reduced. The way in which this is achieved is shown in Fig. 6, which will be explained in more detail below.

The positions and barcode numbers of the markers found in the image are first passed to a process 40, which computes the camera position based on this information using standard photogrammetric techniques. The position and orientation computed, 47, will be liable to change suddenly if a marker disappears or reappears, as explained earlier. The position and orientation are input to a process 41 which computes the positions in the image at which the markers would be expected to have appeared, based on the measured camera position, the internal parameters of the camera, and the positions of the markers in the studio. The difference between the expected and actual marker positions is computed by a subtracter 42. Two subtraction operations are required for each visible marker, one for the horizontal coordinate and one for the vertical coordinate. If the camera calibration and absolute marker locations are perfectly accurate, the output of this subtracter will be zero for each marker; however in a practical situation there will be small differences. Note that the process 40 typically aims to minimise the sum of the squares of these differences. In some implementations of 40, the differences will be generated during the computation process, so that a separate process 41 may not be required. These difference signals (comprising horizontal and vertical displacement differences for each marker) are filtered by a temporal low-pass filter 43. This filter may be, for example, a first-order recursive filter. The filter

may be applied separately to the horizontal and vertical position differences for each marker, although more sophisticated processing that treats both components together is also possible. The filtered position differences are added to the observed marker positions by an adder 44 to form a set of
5 'adjusted' marker positions 50. The position and orientation are computed from these adjusted positions by a process 45, which operates in essentially the same way as process 40. Note that in some implementations of the processes 40 and 45, it is useful to have an estimate of the camera position, for example to initialise a recursive minimisation process. If this is the case, it may be
10 convenient to use the camera position computed by 40 to initialise the process 45. The position computed by 45 may be used to initialise the process 40 on the next frame.

In a steady state with no camera movement and no markers appearing or
15 disappearing, the output of the low-pass filter 43 will be the same as its input, so that the adjusted marker positions 50 will be equal to the expected positions computed by 41. The camera position calculation process 45 will thus compute values for position and orientation that are identical to those computed by 40 and will therefore correspond to those derived directly from the observed
20 marker positions.

If a marker disappears, the camera position and orientation 47 computed from the raw marker data may change slightly. This will produce a change in the expected marker positions computed by 41. However, this change will not
25 immediately appear at the output of the filter 43, due to its temporal low-pass nature. The adjusted marker positions 50 will therefore be essentially unchanged from the adjusted positions corresponding to the situation before the marker disappeared, although there will be one fewer marker positions present. (Note that the subtractor 42 and the adder 44 only produce a valid output for
30 markers which are present at both of their inputs.) The camera position that is computed by process 45 will therefore be essentially unchanged from that which would have been computed had a marker not disappeared. This is

because the marker positions 50 are all self-consistent; that is, any combination of markers will produce the same computed camera position and orientation. As time passes (and assuming no more markers appear or disappear), the output of filter 43 will gradually tend towards its input value.

5 The adjusted marker positions will thus move slowly towards the positions computed by the process 41. Therefore, the final computed camera position 48 will slowly tend towards those at 47, computed from the raw marker data.

10 If a new marker appears in the image, the filter 43 must be initialised to allow that marker to be filtered. That is, values for the previous adjustment amount for the new marker must be assumed, since these are required for the temporal filter to operate. The simplest option would be to assume a previous adjustment amount of zero, but that would be likely to result in a jump in the computed camera position, since the adjustment amount is unlikely to be
15 exactly zero. A preferable option is to derive a value which will ensure that there is no sudden jump in the calculated camera position. A way in which this may be carried out is as follows. When the marker first appears in the image, its presence is ignored by the camera position and orientation computation processes and the temporal filter, which operate only on the markers visible in
20 both the current and preceding image. The expected position of the new marker, based on the computed camera position and orientation, is then computed by process 46. The difference between this expected position and the position at which it actually appeared is computed by subtracter 51 and used to initialise the filter 43. This initialises the filter to the state it would have
25 been in had the new marker always had this difference between expected and observed position. On the next frame, the new marker is used by all the processing stages. Its appearance is likely to result in a change of the camera position and orientation 47 computed from the raw marker data, but this change will not immediately affect the output of the temporal filter 43. The process 45
30 will therefore produce a result very similar to that which it would have produced had the new marker not appeared. Gradually, however, the correction displacements at the output of the filter 43 will tend towards those

corresponding to the camera position computed by 40, and the camera position computed by 45 will therefore tend to that computed by 40.

5 If fewer than three markers are seen in any image, the system will be unable to compute a position or angle, since the number of unknowns (three angles and three position coordinates) outnumber the number of observed values (two image coordinates for three markers). The system may also fail to compute a valid result in some other circumstances, for example if a marker is mis-identified or its measured position is grossly in error. When this happens, it is advantageous to continue to output the last valid measurement until a valid result is obtained. Furthermore, it may be advantageous to initialise the filter 43 to the state it would be in if all markers had zero positional differences. This allows the system to quickly recover from a situation in which it has made some inaccurate measurements.

15

A variant in which markers are positioned in the field of view of the camera will now be described; this may be used to supplement the positional information obtained from the above method, or may be used to obtain positional information independently. This variant employs an optical system based on

20

clusters of markers placed in the scene. Each cluster preferably contains at least three individual markers (for example encoded concentric rings), preferably at differing heights above a reference surface, and more preferably contains at least six individual markers. The clusters are preferably formed as rigid units on which the individual markers are fixed; this enables the relative locations of each marker within each cluster to be known to a high degree of accuracy.

25

In a preferred implementation, as depicted in Fig. 7, each cluster of markers consists of six circles in one shade of blue against a background in a slightly different shade, in order to allow them to be keyed out. The use of two-tone blue for camera position measurement has already been used

30

effectively. The cluster is designed so that the circles are not all co-planar; this allows the camera position to be computed unambiguously from a single cluster, and also facilitates the automatic measurement of the cluster positions. Each marker within the cluster has a barcode to identify it uniquely.

5

The image is analyzed to locate the exact position of each marker within each cluster, from which the camera position and orientation (three position coordinates and three rotation angles) can be computed, given the known positions of the markers in the studio. The internal camera parameters (including focal length and pixel size) need to be known. Where the camera has a zoom lens, these can be determined by sensors on the lens and/or by known methods of detecting zoom and focus settings of a lens from the image. If a fixed lens is used, the settings can be determined by a separate initial calibration procedure, since they remain fixed for a given choice of lens.

10
15

The clusters need not all contain the same number of individual markers, and individual markers may also be disposed at known positions within the image.

20

The only constraint this system places on the positioning of marker clusters is that at least one, and ideally a minimum of around three, need to be visible at all times. As many clusters as needed to cover the area of the action can be used, placed so as not to interfere with the movement of actors. This approach is inherently extendable to very large studios, and allows for scenes involving significant actor and camera movement, such as following actors running down long corridors.

25

Since the clusters can be re-positioned to suit each scene being shot, it is important to include a method to automatically measure their positions to a high accuracy. This may be achieved by capturing a number of images, each showing two or more clusters. The identities of the markers in each image can be determined automatically. A global least squares adjustment

30

determines the relative position and orientation of all markers across the various images.

5 The system is preferably designed to run in real-time using a standard workstation equipped with a frame grabber. To achieve this, it is desirable to track the position of markers from one frame to the next, so that only a small fraction of the image (around the predicted marker positions) needs to be searched. Furthermore, this may allow the use of a simple algorithm to find the centre of each circle. Using such an approach, it has been found possible to
10 measure the camera position at a rate of 25Hz (when tracking 20 circles) using a PC based on a 200MHz Pentium Pro. When the system is started, however, the entire image does need to be searched, which takes several seconds.

15 Although the techniques described are ideally suited to video camera position determination, they can be applied to, for example, determination of position of robot arms or to virtual reality, in which the position of a headset worn by a user is determined.

CLAIMS:

1. A method of determining the position of an object comprising:
providing a plurality of markers at respective reference positions, at least
5 some of the markers being patterned to encode identification information, the
pattern providing a reference point for each marker;
storing information including a measure of the positions of the markers
and information identifying the patterned markers;
obtaining an image of at least a sub-set of said plurality of markers from
10 a camera associated with the object;
processing the image to identify the positions of said markers in the
image and, for each patterned marker in the image, identifying the respective
reference point and decoding said identification information;
determining a measure of the position of the object based on said
15 processing and decoding and based on said stored information.
2. A method according to Claim 1, wherein the pattern comprises a series
of concentric shapes encoding said identification information, the centre of the
concentric shapes of each patterned marker being identified as a reference
20 point for each marker.
3. A method according to any preceding claim, wherein the patterned
markers encode a plurality of bits of information in binary form as a series of
light and dark shapes.
25
4. A method of determining the position of an object comprising:
providing a plurality of markers at respective reference positions, at least
some of the markers being patterned to encode identification information in
binary form as a series of light and dark regions, the marker also providing a
30 detectable reference point;
storing information including a measure of the positions of the markers
and information identifying the patterned markers;

obtaining an image of at least a sub-set of said plurality of markers from camera means, preferably monochrome camera means, associated with the object;

5 processing the image to identify the positions of said markers in the image and, for each patterned marker in the image, decoding said identification information;

 determining a measure of the position of the object based on said processing and decoding and based on said stored information.

10 5. A method of producing a marker encoding a plurality of bits of information, the method comprising obtaining the information to be encoded in binary form and providing on a backing a series of concentric shapes, each shape having either a high or a low brightness or reflectivity, the thickness or brightness or reflectivity of each ring shape selected according to a pre-
15 determined coding scheme to encode the information.

 6. A method according to any of Claims 2 to 5, wherein the shapes are circular rings.

20 7. A method according to any of Claims 3 to 6, wherein 3 or more bits of identification information are encoded.

 8. A method according to any of Claims 3 to 7, wherein 4 or more shapes are provided.
25

 9. A method according to any of Claims 3 to 8, wherein each bit of the information is assigned to a corresponding shape.

30 10. A method according to any of Claims 3 to 9, wherein at least one outer region of a predetermined brightness is provided.

 11. A method according to any of Claims 3 to 10, wherein at least one parity

or other error-correcting bit is included in the encoded information.

12. A method according to any of Claims 3 to 11, wherein dark shapes are substantially black.

5

13. A method according to any of Claims 3 to 12, wherein light shapes are retro-reflective.

10

14. A method according to Claim 13, further including illuminating the markers with a light source substantially coincident with the axis of said camera means.

15

15. A method according to Claim 2, wherein each shape is coloured, patterned or shaded to encode a plurality of bits of identification information.

16. A method according to any preceding claim, wherein all markers are patterned to encode identification information.

20

17. A method according to any preceding claim, wherein at least some markers encode identification information uniquely identifying that marker among the plurality of markers.

25

18. A method according to any preceding claim, wherein the markers are so disposed that the position of the object can be uniquely determined at all points within a given volume.

30

19. A method according to any preceding claim including adjusting the determined position to smooth out sudden changes in determined position as markers are revealed or obscured.

20. A method according to Claim 19, wherein adjusting comprises applying correction factors to the stored or measured positions of the markers so that

determination of object position is based on corrected marker positions, the correction factors being calculated so that the corrected marker positions tend to mutually self-consistent values, the rate of variation of the correction factors being kept below a pre-determined level.

5

21. A method of determining the position of an object based on identification of a plurality of markers at respective reference positions, the method comprising applying correction factors to stored or measured marker positions to produce respective corrected marker positions and determining the object position based on the corrected marker positions, wherein the correction factors are calculated so that the corrected marker positions tend to mutually self-consistent values, the rate of variation of the correction factors being limited to a pre-determined level.

10

22. A method according to Claim 20 or 21, wherein the pre-determined level is selected so that transitions in the determined image position resulting from one or more identified markers becoming unusable or new markers becoming usable are smoothed over a period of at least about 2 seconds.

15

23. A method according to any preceding claim, wherein a measure of the accuracy of position determination is obtained.

20

24. A method according to Claim 23 as dependent on any of Claims 20 to 22, wherein said measure is based on determination of said correction factors.

25

25. A method according to any preceding claim, wherein the markers are positioned at varying distances from a reference surface, the stored information including a measure of the distance of the marker from the reference surface.

30

26. A method according to any preceding claim wherein said object is a video camera having a field of view differing from the field of view of said camera means.

27. A method according to Claim 26 further comprising providing at least one supplementary marker at a given position so as to be viewable by the video camera, the or each marker having identification information encoded thereon in binary form as a series of concentric rings of two tones, the tones being
5 selected to enable the marker to be keyed out of the video camera image by chroma-keying, said determining the position of the video camera being based on the position of the supplementary marker, if present, in the video camera image.

10 28. A method of determining or correcting a measure of the position of a video camera the method comprising providing at least one marker at a given position viewable by the video camera, the or each marker having identification information encoded thereon in binary form as a series of concentric rings of two tones, the tones being selected to enable the marker to be keyed out of the
15 video camera image by chroma-keying, a measure of the position of the video camera being determined based on the position of the marker, if present, in the video camera image.

20 29. Apparatus for position determination comprising a camera for mounting on an object whose position is to be determined; memory means arranged to store measures of the positions of a plurality of markers and to store information identifying patterned encoded markers; image processing means arranged to process an image output by the camera to identify the positions of
25 markers in the image; decoding means for decoding information encoded in the patterned markers; and position determining means arranged to determine the position of the object based on the output of the image processing means, the decoding means and the information stored in the memory means.

30 30. Apparatus according to Claim 29 wherein the memory means, image processing means, decoding means and position determining means are integrated into a single computer.

31. Apparatus according to Claim 29 or 30 further including a plurality of markers, including a set of patterned markers having identification information encoded as a series of concentric rings of light and dark regions.

5 32. Apparatus according to any of Claims 29 to 31 for use with markers including retro-reflective portions, the apparatus having light-projection means substantially optically coincident with the axis of said camera means.

10 33. Apparatus according to Claim 32 wherein the light-projection means is pulsed, arranged to be pulsed on at a time when said camera means is integrating.

15 34. Apparatus according to Claim 32 or 33 wherein the camera means is provided with filter means having a transmission window selected to pass light transmitted by the light-projection means but to inhibit at least some other light.

20 35. Apparatus according to Claim 34, wherein the light-projection means comprises an infra-red LED and said filter inhibits visible light but passes infra-red.

25 36. Apparatus according to any of Claims 29 to 35 including means for applying correction factors to stored or measured marker positions to produce respective corrected marker positions and determining the object position based on the corrected marker positions, wherein the correction factors are calculated so that the corrected marker positions tend to mutually self-consistent values, the rate of variation of the correction factors being limited to a pre-determined level.

30 37. Apparatus according to any of Claims 29 to 36 arranged to determine the position of a video camera wherein the camera means comprises a further camera mounted on the video camera, preferably facing in a different direction to the video camera.

38. Apparatus according to Claim 37, including means for detecting the zoom and/or focus settings of the video camera.

5 39. A set of patterned markers each comprising a series of concentric rings, each marker encoding a plurality of bits of information in binary form as the brightness or reflectivity and/or thickness of each ring, the information encoded varying between the markers.

10 40. Use of a marker comprising series of concentric rings to encode a plurality of bits of information in binary form as the brightness or reflectivity and/or thickness of each ring, the centre of the rings providing a reference point to enable the position of the marker to be determined.

15 41. A method according to Claim 28, wherein a plurality of said markers are first mounted on a base to form a cluster of markers for providing at a position viewable by said video camera.

20 42. A method according to Claim 42, wherein at least three non-co-planar markers are mounted on said base.

25 43. Marker apparatus for positioning in the field of view of a video camera for use in determining the relative position of the camera and the marker apparatus, the apparatus comprising at least three non-co-planar patterned markers mounted at predetermined relative positions on a substantially rigid base, the patterned markers each having identification information encoded thereon in binary form.

30 44. Apparatus according to Claim 43, wherein the information is encoded as a series of concentric rings of two tones, the tones being selected to enable the marker to be keyed out of the video camera image by chroma-keying.

45. A method of determining the relative positions of at least one camera and an object having a plurality of markers, preferably at least three, mounted thereon, the relative positions of the markers and the object preferably being substantially fixed, each marker being patterned to encode identification information in binary form, preferably as series of concentric rings, the method comprising storing the relative positions of the markers and information enabling the identification information to be decoded, identifying the positions of the markers and decoding the identification information encoded thereon, and, based on the stored information, determining a measure of the relative positions and orientations of the or each camera and the object.
46. A method substantially as any one herein described, with reference to the accompanying drawings.
47. Position determination apparatus substantially as herein described, with reference to the accompanying drawings.

1 / 6




Fig. 1

2/6




Fig. 2

3/6

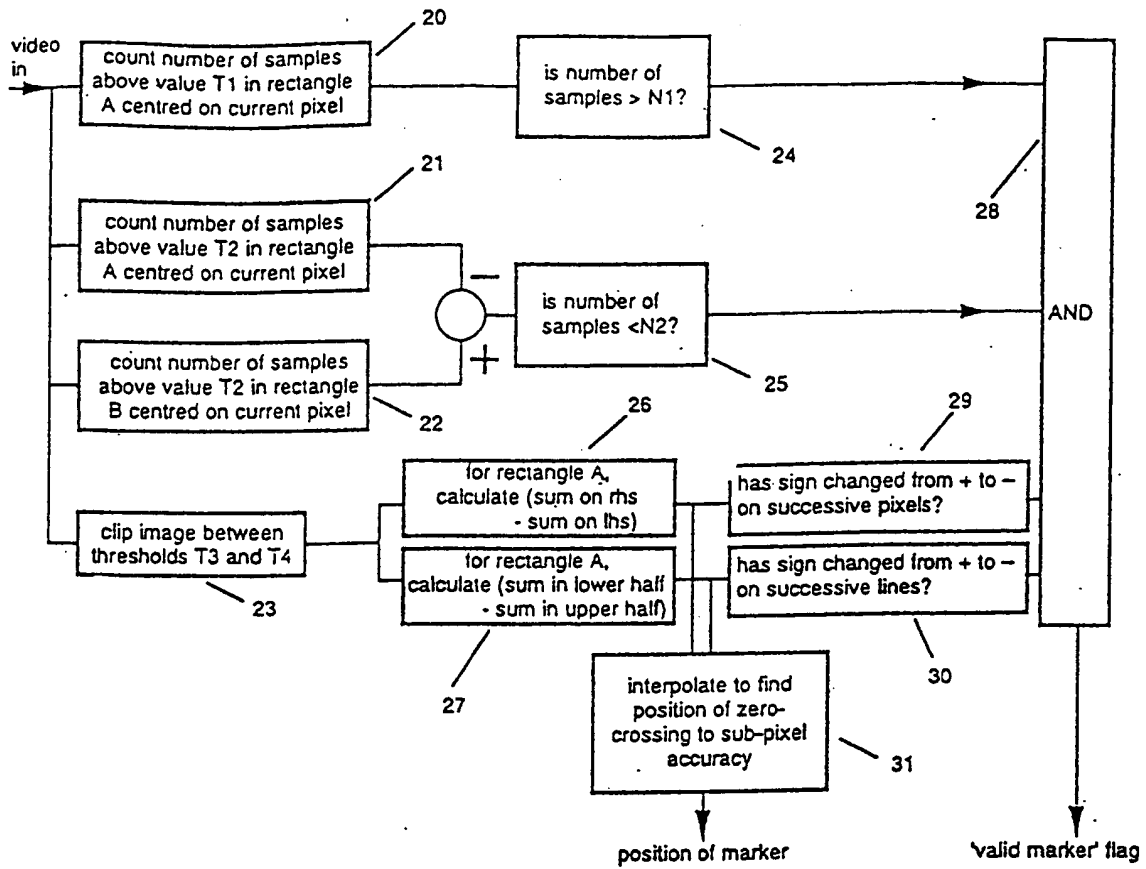


Fig. 3




Fig. 4

4/6

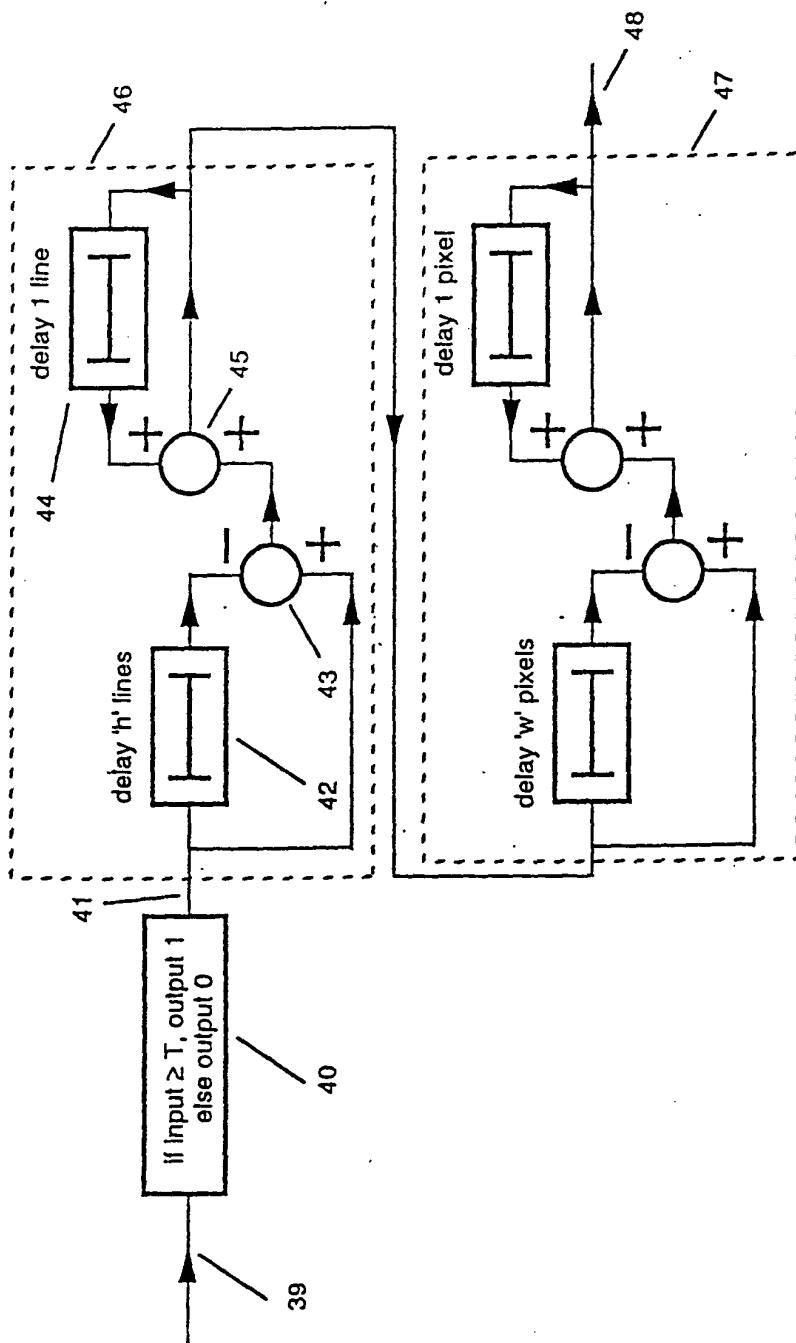


Fig. 5




Fig. 6

6/6




Fig. 7

INTERNATIONAL SEARCH REPORT

Inter national Application No

PCT/GB 98/01576

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G01S5/16 G01S17/87 H04N5/272

According to International Patent Classification(IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G01S H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 97 11386 A (OMNIPLANAR INC) 27 March 1997	1,3,4,8, 10, 12-21, 25,26, 29-32, 35-37, 39-43
Y	see page 2, line 35 - page 10, line 8 see page 10, line 33 - page 11, line 7 see page 12, line 25 - line 37 --- -/--	2,5,6,9, 27,28, 33,34, 44,45

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

9 September 1998

Date of mailing of the international search report

16/09/1998

Name and mailing address of the ISA
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2220 HV Rijswijk

Authorized officer

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 98/01576

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 706 105 A (CONSORZIO TELEROBOT) 10 April 1996	1,4,7, 10,12, 15-18, 20,21, 29-31, 36,39,40 2,5,6,9
Y	see column 2, line 18 - column 3, line 27 see column 8, line 30 - column 9, line 19 ---	
Y	US 5 227 985 A (DEMENTHON DANIEL F) 13 July 1993 see column 5, line 8 - column 6, line 29 ---	45
Y	GB 2 305 050 A (ORAD HI TEC SYSTEMS LTD) 26 March 1997 see page 8, line 19 - page 14, line 12 ---	27,28,44
Y	EP 0 265 542 A (RATHBONE RICHARD R ; VALLEY ROBERT A JR (US); KINDLMAN PETER J (US)) 4 May 1988 see column 3, line 8 - column 4, line 45 -----	33,34

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 98/01576

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9711386 A	27-03-1997	EP 0852732 A	15-07-1998
EP 0706105 A	10-04-1996	IT 1271241 B	27-05-1997
US 5227985 A	13-07-1993	NONE	
GB 2305050 A	26-03-1997	AU 6935396 A	27-03-1997
		EP 0848886 A	24-06-1998
		WO 9709830 A	13-03-1997
		NO 981010 A	09-03-1998
EP 0265542 A	04-05-1988	NONE	